Perfect and Maximum Randomness in Stratified Sampling over Joins

Niranjan Kamat Arnab Nandi
Computer Science & Engineering
The Ohio State University
{kamatn,arnab}@cse.osu.edu

arXiv:1601.05118v1 [cs.DB] 19 Jan 2016

ABSTRACT

Supporting sampling in the presence of joins is an important problem in data analysis. Pushing down the sampling operator through both sides of the join is inherently challenging due to data skew and correlation issues between output tuples. Joining simple random samples of base relations typically leads to results that are non-random. Current solutions to this problem perform biased sampling of one (and not both) of the base relations to obtain a simple random sample. These techniques are not always practical since they may result in the sample size being greater than the size of the relations due to sample inflation, rendering sampling counter-productive.

This paper presents a unified strategy towards sampling over joins, comprising two key contributions. First, in the case that perfect sampling is a requirement, we introduce techniques to generate a *perfect* random sample from both sides of a join. We show that the challenges faced in sampling over joins are ameliorated in the context of stratified random sampling as opposed to simple random sampling. We reduce the dependency of feasibility of sampling from relation level to strata level. Our technique minimizes the sample size while maintaining perfect randomness.

Second, in the case that random sampling is not a requirement but is still preferred, we provide a novel sampling heuristic to *maximize* randomness of the join. It allows us to allocate a fixed sample size between multiple relations consisting of multiple strata to maximize the join randomness. We validate our techniques theoretically and empirically using synthetic datasets and a standard benchmark.

1. INTRODUCTION

Data sampling has greatly increased in popularity, as evidenced by inclusion of the `TABLESAMPLE` keyword into the `SQL` standard, with implementations in most mainstream databases. In cases where the growth of data volume outpaces the hardware or software's ability to process it in time, sampling presents a pragmatic approach towards providing insights at scale. Incorporating sampling into analytics not only speeds up the analysis cycle but also potentially reduces the resource requirements by orders of magnitude.

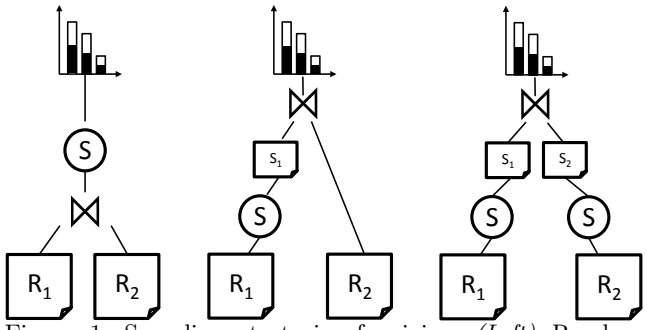

Figure 1: Sampling strategies for joins: *(Left)* Random Sampling after the join execution provides correct semantics, but the entire relations have to participate in the join. *(Middle)* Works by Chaudhari et al. [6] and Gibbons et al. [10] have provided means for sampling from one side of the join; this poses several challenges if both the relations are large. *(Right)* Our first contribution is to provide techniques to sample from both sides of the join to minimize the total sample size while maintaining perfect join randomness.

Due to the non-exact nature of results, the quality and method of sampling is a key consideration. A random probability sample as compared to a non-probability sample is usually necessary in order to obtain meaningful results. While the use of non-random samples has been considered in aggregation queries [15], the use of random samples allows the user to utilize statistical methods over the data in a principled manner. The nature of the sample also impacts visual representations – a visualization of a sample may look very different from the original if it is not sampled correctly.

Techniques such as simple random sampling and stratified random sampling can be used sample the data. Stratified random sampling is usually preferred since it provides a better representation of the data, yielding higher efficiency (lower variance) of the measure [7]. Stratified sampling is especially suited for joins, since the data can be considered to be stratified by the groups in the join column.

In the context of `JOIN` queries, sampling is a far more compelling, yet challenging task. It is generally preferred for the sampling techniques to preserve the randomness of the join output. Performing a join can be expensive – in a join between relations R_1 and R_2, the minimum time complexity of using a simple hash join would be $O(|R_1| + |R_2|)$, and

$O(|R_1| \times |R_2|)$ for a simple nested loop join. Therefore, performing sampling after the join result is materialized, may not be a pragmatic approach.

Another alternative approach is to produce the sampled join result by sampling the base relations themselves. Consider the potential rewrite:

`TABLESAMPLE(A ⋈ B)`
to
`TABLESAMPLE(A) ⋈ TABLESAMPLE(B)`

where ⋈ denotes a new join algorithm. If possible, this approach would have significant benefits – in the case of visual data exploration tools over large dataset, a visualization of the join result can be approximated from samples available in the local cache. In the case of distributed join execution, the sampling of the base relations can be performed at the *Map* step, allowing for the entire join to be performed as a single MapReduce job.

However, obtaining random samples of a join by sampling base relations (also termed *join sampling*) is an inherently difficult problem from a correctness perspective [1, 6, 24]. Common solutions include sampling from only one side of the join, or presenting a non-representative sample to the user. In order to highlight the challenges involved, we now walk through an example to motivate our line of inquiry.

1.1 Motivating Example

Excellent examples have been provided by Gibbons et al. [1] and Chaudhari et al. [6] that demonstrate the inherent difficulty of join sampling. We briefly revisit the example given in the latter, since it motivates our algorithms.

Consider joining two relations R_1 and R_2 with the values in the join column being $\{1, 2...2\}$ for R_1 and $\{2, 1...1\}$ for R_2, such that there are 99 tuples that have the value 2 in R_1 and have the value 1 in R_2. In a simple random sample of the join between R_1 and R_2, on an average, half the tuples would be joined on 1 and the other half on 2. However, it is clear that if both relations are sampled randomly in an unweighted fashion, their resultant join would more likely be an empty set. This motivates us towards using weighted sampling and guides us towards the reasons behind simple random sampling of joins being difficult.

1.2 Why Stratified Random Sampling?

Using the motivating example, we enumerate the difficulties involved in obtaining *simple* random samples of joins. We then show how they can be alleviated using *stratified* random sampling.

Sample Inflation: The key reason behind infeasibility of using simple random sampling in joins is sample inflation. In order to avoid correlation in output tuples, a sampled tuple can be joined only once with another tuple. Therefore, it is possible for the size of the samples to exceed size of the input relations, resulting in sampling being counter-productive.

In stratified sampling, the data is divided into strata and simple random sampling is then applied within each stratum. Thus, the different strata can be sampled independently of one another. If the sample size needed for a stratum exceeds the stratum size, we can choose to not sample it. The entire base stratum can be added to the sample. While performing the join, tuples can be chosen from the stratum randomly. In this way, we are able to handle sample inflation and minimize the sample size. A detailed explanation is provided in Section 4.4.

Smaller Output Size & Sample Tuple Loss: In sparse or skewed data, using simple random sampling, it is possible for the join of the samples to be smaller than intended. If both the relations involved in the join are sampled without a smart sampling strategy, there may be an absence of joinable tuples on both sides, yielding a smaller output size than expected. This also renders some of the sampled tuples non-productive.

This challenge can be handled by using stratified random sampling by choosing equal number of tuples from both relations for each stratum. This makes sure that all sample tuples are utilized and the sample size is as expected.

Non-representative Output: Using simple random sampling, while samples will be random, and representative in the average case, for any non-trivial sampling fraction (lesser than 1), there is a non-zero likelihood of some groups being left out of the join. This issue is resolved using stratified random sampling since join will contain tuples from all strata.

1.3 Overview & Contributions

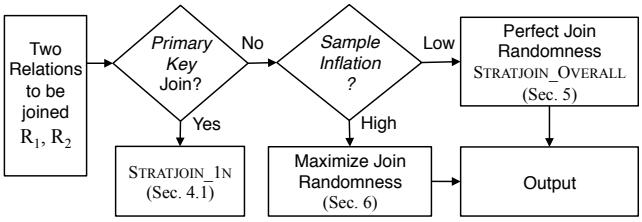

Figure 2: Overview

Figure 2 presents an overview of our *perfect* and *maximum* randomness solutions to the join sampling problem. Sampling in the presence of primary key and foreign key relationship has been investigated by the *Aqua* system, StratJoin_1n (Section 4.1) presents a stratified version of their join algorithm. If the join has a low level of sample inflation (Section 4.4), we can use StratJoin_Overall (Section 5) to obtain a *perfectly* random output. If the reduction in sample size is low due to sampling inflation, we provide techniques to *maximize* the sample randomness and use a correlated output sample (Section 6). The decision to use a perfectly random or maximally random sample is ultimately user-dependent.

We now describe the two main contributions of this paper in greater detail. First, we provide means to obtain stratified random samples of equi-joins using samples of either one or both relations. The use of stratified random sampling alleviates some of the issues involved in simple random sampling. We introduce a generalized algorithm, StratJoin_Overall, that gives us control over deciding the ideal sampling as well as join strategies across different strata independently of each other.

Our second contribution is in the case where random samples are not mandated but preferred nonetheless. Here, a novel algorithm is presented that maximizes the *randomness* of a sample given the constraint of fixed sample size. Towards this end, it provides a strategy for allocating samples to the different strata of the different relations involved in the join. These results have been proven theoretically as well as validated experimentally using well defined metrics.

2. RELATED WORK

Join sampling has been incorporated in systems such as Microsoft SQL Server, IBM DB2 and more recently BlinkDB. While these approaches involve techniques implemented at various layers of the database such as table scans, offline catalogs, and aggregation, there still exist several open opportunities. In order to obtain random samples of join, the typical approach is to sample a single relation [2,3,6,24] and use the entire second relation in the join.

Work by Olken et al. shows that it is easy to commute selection with sampling but projections and joins are harder [24]. They sample a single relation.

Chaudhari et al. provide better algorithms for joins using sample of a single relation [6] and consider different availabilities of statistics and indexes. They describe reasons behind the inherent difficulty of obtaining a simple random sample of join while extending sampling towards multiple relations.

The *Aqua* system [1] provides techniques to obtain a simple random sample of the join involving relations having primary key and foreign key relationships. This is an easier case to handle since sampled tuples do not result in correlation between the output tuples.

Research in obtaining random samples of joins has stalled since inherent problems involved in simple random sampling of joins were presented [6]. However, a lot of effort has been devoted towards finding meaningful aggregate estimates in presence of correlation. Online aggregation and ripple joins provided join evaluation techniques to obtain meaningful error estimation during query execution [11–15, 22].

In the SciBORQ system, Sidirourgos et al. [25] use join sampling to maintain correlation between the join attributes in order to provide the query results with greater precision. It plays an important role in their design of impressions, which consist of multi-level, workload-adaptive query-focused samples.

Jermaine et al. remove the dependence of the ripple join algorithms on the data residing in memory in order to provide statistical guarantees for the estimates [16, 17]. The DBO [18] and Turbo-DBO [9] systems elevated the approach towards processing multiple relations in a scalable fashion. Nirkhiwale et al. present the sampling algebra inherent in these techniques [23].

3. PRELIMINARIES

In this section, we introduce preliminaries and the terminology used in this paper. We enumerate the criteria that need to be satisfied for samples to be termed either as simple random samples or stratified random samples. We use these criteria to prove the correctness of the algorithms [7]. Next, we define a join operator which is used in our algorithms (*Mini-Join*). We provide a simple extension to *Stream-Sample* to sample from both sides (Srs_Both). Our notations resemble those used by Chaudhari et al. [6] and are shown in Table 1.

Join Sample Size Definition: Total sample size going in to join is defined as the sum of sizes of samples when first joined output tuple is emitted. This definition is intuitive and independent of the implementation.

3.1 Random Sampling Criteria

We can show that a sample is a simple random sample or a stratified random sample by showing that they satisfy the criteria given below respectively.

Symbol	Explanation
R_i	i^{th} Base Relation
N_i	Cardinality of R_i
S_i	Sample of R_i
n_i	Cardinality of S_i
f	Join Sampling Rate
A	Join Column
a	Join Value
$t.A$	Value of Column A in Tuple t
$m_i(a)$	Number of Tuples in R_i with Value a in A
m_i^j	Number of Tuples in R_i belonging to j^{th} stratum
$mm_i(a)$	Number of Tuples in S_i with Value a in A
mm_i^j	Number of Tuples in S_i belonging to j^{th} stratum
S_{ij}	Sampling-based join between R_i and R_j
S_{ij}^k	Tuples in k^{th} stratum of S_{ij}

Table 1: List of Notations

3.1.1 Simple Random Sampling Criteria

We briefly present the two criteria that need to be satisfied for a sample to be termed as a simple random sample [7].

***SRS_1*:** All tuples in the population should have the same probability of occurrence.

***SRS_2*:** All samples of a given size should be equally likely, i.e., a tuple being present in the sample should not affect the probability of another.

3.1.2 Stratified Random Sampling Criteria

In stratified random sampling, the population is first divided into strata such that the strata are mutually exclusive and the union of the strata covers the entire population. Samples are then chosen from each stratum using simple random sampling. The criteria can be given as follows.

***StRS_1*:** Every sample stratum should have the intended size depending on the allocation type (proportional, optimal, cost-based, etc.).

***StRS_2*:** All tuples belonging to a stratum should have the same probability of occurrence.

***StRS_3*:** Within a stratum, a tuple being present in the sample should not affect the probability of another.

These criteria clearly show the ease of obtaining stratified random samples over simple random samples since randomness only needs to be achieved at the stratum level.

3.2 The Mini-Join Operator

If all tuples within a sample stratum are joined with all tuples in the corresponding stratum of the other, the output tuples will be correlated. The *Mini-Join* operator, denoted by $\tilde{\bowtie}$, is designed to avoid this correlation. Instead of producing all $|mm_1(a) \times mm_2(a)|$ tuples of the cartesian product for a stratum a, random tuples from both samples are selected and joined. A sample tuple is used only in a single output tuple. Each stratum sample will be a sample without replacement of the non-sampled join between the strata. The complexity of this operator is $\sum_{a \in Strata} min(mm_1(a), mm_2(a))$, where $Strata$ is the set of all strata in $S_1 \bowtie S_2$. This operator is influenced by similar actions recurring in [6].

Algorithm:

For every stratum $a \in Strata$:
  Repeat $min(m_1(a), m_2(a))$ times:
    Choose hitherto non-selected random tuples from $m_1(a)$ and $m_2(a)$, join them, and output the resulting tuple.

Example:
Consider two relations given below that are specified as $(stratum_value : stratum_count)$:
$R_1 = (a : 2, b : 5)$ and $R_2 = (a : 3, b : 3)$.
The output counts after applying the *Mini-Join* operator will be $R_1 \tilde{\bowtie} R_2 = (a : 2, b : 3)$. Thus, projecting the joined tuples on any relation will result in distinct tuples. For example, for the first stratum, for each tuple from R_1 we will select a different tuple from R_2.

3.3 Simple Random Sampling Over Joins

In this section, we review the existing techniques used to obtain simple random samples over joins by sampling a single relation – *Simplified Aqua* and *Stream-Sample*. We provide a simple extension to *Stream-Sample* that enables sampling both sides of the join to obtain a simple random sample over join – SRS_BOTH. Finally, we provide the distribution of strata counts in a simple random sample over joins to better aid understanding of join sampling.

3.3.1 Simplified AQUA Algorithm

We briefly describe a simplified version of the basic join sampling algorithm used in the AQUA system [1] as we provide a stratified version of it in Section 4.1.

Algorithm:
1. Perform simple random sampling over the foreign key table.
2. Join the sample with the primary key table.

Theorem 1 (Based on Gibbons et al. [10] Theorem 4.3) *Simplified Aqua* results in a without-replacement simple random sample of $R_1 \bowtie R_2$, while sampling a *single* relation.

3.3.2 Strategy Stream-Sample

We briefly describe the *Strategy Stream-Sample* algorithm given by Chaudhari et al. [6], since our algorithms are based on the same problem formulation. The *Stream-Sample* algorithm is designed for obtaining a simple random sample of the equi-join between R_1 and R_2 by sampling a single relation for the general case of multiple tuples from a relation joining with multiple tuples from another.

Algorithm:
1. Obtain a with-replacement sample S_1 of R_1, where the sampling weight $w(t_1)$ for a tuple t_1 is set to $m_2(t_1.A)$.
2. For each tuple from S_1 in a streaming fashion do:
   (a) Sample a random tuple t_2 from amongst all tuples $t \in R_2$ that satisfy $t.A = t_1.A$.
   (b) Output $t_1 \bowtie t_2$.

Theorem 2 (Based on Chaudhari et al. [6] Theorem 6) *Stream-Sample* gives a with-replacement sample of $R_1 \bowtie R_2$, while sampling a *single* relation.

Discussion:
The proof of correctness is provided in Section 10.1. Since the sample size is $f \times |R_1 \bowtie R_2|$ and hash-index is available on R_2, the join complexity will be $f \times |R_1 \bowtie R_2|$. *Group-Sample* and *Frequency-Partition-Sample* deal with the availability of indexes and statistics [6]. In this paper, we have considered the *Stream-Sample* algorithm because it is the simplest amongst them. It helps illustrate the basic ideas behind join sampling clearly. Further, the sample size inflation issue which hurts *Stream-Sample* affects the other two algorithms as well.

3.3.3 SRS_BOTH

We provide a method to sample both sides of a join in order to obtain a simple random sample of the join by making a *minor* modification to *Stream-Sample*. If the relations are sampled independently, the number of output tuples, using the *Mini-Join* operator, will be $\sum_{a \in Strata} min(mm_1(a), mm_2(a))$. Therefore, in order to avoid loss of sampled tuples, after sampling a relation, the cardinality information of each sampled stratum should be used to sample the other relation. It is easy to notice that *Stream-Sample* can be modified so that instead of selecting and joining tuples from R_2 while tuples in S_1 are streaming by, the selected tuples are added to S_2. Once the entire S_2 is materialized, tuples that would have joined in *Stream-Sample* are joined now instead.

Algorithm:
1. Obtain a with-replacement sample S_1 of R_1, where the sampling weight $w(t_1)$ for a tuple t_1 is set to $m_2(t_1.A)$.
2. For each tuple from S_1 in a streaming fashion do:
   (a) Sample a random tuple t_2 from amongst all tuples $t \in R_2$ that satisfy $t.A = t_1.A$.
   (b) Add t_2 to the sample S_2 while preserving the insertion order of tuples.
3. For each tuple from S_1 in a streaming fashion, in the same order as before, do:
   (a) Choose the next tuple t_2 from S_2, and output $t_1 \bowtie t_2$.

Theorem 3 SRS_BOTH results in a with-replacement simple random sample of $R_1 \bowtie R_2$, while sampling *both* relations.

Proof of Correctness:
SRS_BOTH can be considered as a *minor* modification of *Stream-Sample*, where instead of choosing tuples from R_2 as the tuples from S_1 are streaming by, the chosen tuples from R_2 are used to materialize S_2 and joined later. Both algorithms result in the same join output. A detailed proof of correctness for *Stream-Sample* in Appendix 10.1.

3.3.4 Strata Counts in SRS over Joins

The strata counts follow a multinomial distribution with $p_i = \frac{m_1^i \times m_2^i}{|R_1 \bowtie R_2|}$, where p_i is the probability of any chosen sample tuple belonging to stratum i. Thus, the number of tuples in the output belonging to a stratum might vary in successive trials. It is possible for smaller strata to not have any tuples in the output. Stratified random sampling based algorithms resolve these issues.

4. STRATIFIED RANDOM SAMPLING OVER JOINS

The join sampling problem for obtaining perfectly random samples can be formally defined as follows:

Given relations to be joined, obtain a stratified random sample of the overall join by sampling one or both relations.

First, we provide techniques for sampling a single side of the join (Sections 4.1 and 4.2). Then, we extend our methods towards sampling both the sides of the join (Section 4.3). We note that in order to avoid correlated samples in the output, sampling in the general case of non-primary key based join can become expensive (Section 4.4). We provide means to alleviate this problem as well. The combination of these methods forms our overall algorithm (Section 5).

It should be noted that our contribution is towards the overall algorithm and its correctness; efficient implementations, performance optimizations, and index constraints are

beyond the scope of this paper. The algorithms are detailed for a single node configuration. They use hash-based indexes in order to simplify the sampling and join complexities, since the index type is orthogonal to our algorithms' correctness. We assume availability of indexes on the join columns of all relations in order to speed up both sampling and join.

4.1 StratJoin_1n

Aqua considers the important use case of joins based on primary key and foreign key relationship and is easier to handle. A tuple in R_1 can join with a single tuple from R_2. The sample size required to achieve a sampling rate is smaller than in the general join case due to not having to avoid sample correlation. We obtain the benefits of sampling without incurring any of the additional costs due to sample size inflation. It is an important use case as shown by the fact that 13 of the 17 queries in the TPC-D benchmark have such joins [1]. Setting the correct strata sizes, we modify the join algorithm used in the *Aqua* system as given below.

Algorithm:
1. From each stratum in base relation R_1, sample f fraction of the tuples without replacement giving S_1.
2. Join S_1 with R_2.

Complexity:
Sampling complexity is equal to the size of S_1, $n_1 = f \times |R_1|$. Since each tuple in S_1 will join with a single tuple from R_2, join complexity will equal n_1.

Theorem 4 StratJoin_1n results in a without replacement *stratified* random sample of $R_1 \bowtie R_2$, while sampling a *single* relation.

Proof of Correctness:
Satisfaction of $StRS_1$: A stratum a has the correct number of tuples, $f \times m_1(a)$, since in the absence of sampling, it would have $m_1(a)$ tuples.
Satisfaction of $StRS_2$ & $StRS_3$: Since each stratum in S_1 is a sample without replacement and will join with a single tuple from R_2, each output stratum will be a sample without replacement and will therefore satisfy $StRS_2$ & $StRS_3$.

Discussion:
Note that R_2 is usually smaller (dimension table) than R_1 (fact table). Further, sampling from R_2 wastes tuples that are sampled from both relations which do not join. Therefore, we can see that sampling from both relations will be counter-productive. StratJoin_1n can be extended to join multiple foreign key based tables by simply performing a normal join of the other tables with $S_1 \bowtie R_2$.

4.2 StratJoin_nn

We now provide an algorithm for obtaining stratified random samples of joins by sampling a single relation, based on *Stream-Sample*. This algorithm handles the general case where multiple tuples from a relation can join with multiple tuples from another. We will then build upon these concepts to design algorithms to sample both sides (StratJoin_Both and StratJoin_Overall).

Algorithm:
1. For a stratum a, use sampling with replacement to sample $n(a) = f \times m_1(a) \times m_2(a)$ tuples from $m_1(a)$ tuples.
2. For each sample tuple, join with a randomly chosen tuple from amongst the $m_2(a)$ tuples of R_2 that can join with it.
3. Output the joined tuple.

Complexity:
The sampling complexity is determined by the size of the sample S_1, $O(f \times |R_1 \bowtie R_2|)$. Each tuple in S_1 is joined with a single tuple from R_2 chosen randomly and, thus, the join complexity will be defined by the number of join operations, $O(f \times |R_1 \bowtie R_2|)$.

Theorem 5 StratJoin_nn results in a with-replacement *stratified* random sample of $R_1 \bowtie R_2$, while sampling a *single* relation.

Proof of Correctness:
Satisfaction of $StRS_1$: A stratum a has the correct number of tuples, $f \times m_1(a) \times m_2(a)$.
Satisfaction of $StRS_2$ & $StRS_3$: Any stratum a in S_1 is a sample with replacement from the corresponding stratum in R_1.Thus, every tuple in a sample stratum is equally likely as every other tuple in the same sample stratum. Similarly, while selecting a tuple from R_2 to join with it, each tuple amongst the $m_2(a)$ tuples is equally likely, since it is randomly chosen. Therefore, for every joined output tuple, every tuple amongst the $m_1(a) \times m_2(a)$ tuples is equally likely, satisfying $StRS_2$. Further, selecting a tuple in S_1 does not influence the selection of another tuple in the same sampled stratum and similarly, since tuples from R_2 are chosen randomly to join with it, selecting a tuple from R_2 does not affect selection of another, thus satisfying $StRS_3$. Thus, every sampled stratum is a sample with replacement from the corresponding non-sampled stratum.

Discussion:
We note that the StratJoin_nn algorithm is a stratified version of *Stream-Sample* and suffers from its sample inflation issue. Techniques given in Section 5 can be used to improve StratJoin_nn by handling sample inflation. Note that BlinkDB [3] supports stratified random sampling over a single relation and does not consider the both-sides approach discussed in this paper and does not handle sample inflation either.

4.3 StratJoin_Both

The central idea used in sampling a single relation presented in StratJoin_nn extends easily for sampling both relations, by sampling $n(a) = f \times m_1(a) \times m_2(a)$ tuples from both the relations for a stratum a. This algorithm handles the general case of sampling both sides of a join.

Algorithm:
1. Generate samples S_1 and S_2 as follows:
   (a) For each stratum a in R_1, select $f \times m_1(a) \times m_2(a)$ tuples from amongst $m_1(a)$ tuples, using sampling with replacement giving S_1.
   (b) For each stratum a in R_2, select $f \times m_1(a) \times m_2(a)$ tuples from amongst $m_2(a)$ tuples, using sampling with replacement giving S_2.
2. Apply the *Mini-Join* operator between S_1 and S_2.

Complexity:
The complexity of constructing the samples is equal to the sample size, $O(f \times |R_1 \bowtie R_2|)$. Since the *Mini-Join* operator is used to join the two samples, and the two samples have identical sizes in every stratum, the join complexity is equal to $O(f \times |R_1 \bowtie R_2|)$.

Theorem 6 StratJoin_Both results in a *stratified* random sample with replacement of $R_1 \bowtie R_2$, while sampling *both* the relations.

Proof of Correctness:
Satisfaction of $StRS_1$: A stratum a has the correct number of tuples, $f \times m_1(a) \times m_2(a)$, since in the absence of sampling, it would have consisted of $m_1(a) \times m_2(a)$ tuples.
Satisfaction of $StRS_2$ & $StRS_3$: Each sampled stratum is a sample with replacement of the corresponding relation stratum. Selecting tuples from each sample stratum to join using the *Mini-Join* operator does not introduce any correlation (Section 3.2). Thus, output for stratum a is a sample with replacement from amongst $m_1(a) \times m_2(a)$ tuples.

4.4 Avoiding Sample Inflation

The algorithms presented so far result in stratified random samples of the joins. However, the non-primary key based algorithms still suffer from sample inflation due to avoidance of correlation amongst output tuples. The techniques used in this section and the next can be used to rectify this issue.

It is well-known that in order to obtain simple random samples over joins, the sampling rates f_1 and f_2 over relations R_1 and R_2 respectively need to be chosen such that $f_1 \times f_2 \geq f$ [6]. In our algorithms StratJoin_Both and StratJoin_nn, since we work at the level of strata, the sampling fraction of a stratum a in R_1 can be given by $f_1(a) = \frac{f \times m_1(a) \times m_2(a)}{m_1(a)} = f \times m_2(a)$. This constraint is more flexible since the sampling fraction is dependent only on the size of the corresponding stratum in the other relation.

Even so, the sampling fraction can still be greater than 1, which is counter-productive. Therefore, depending on the strata sizes and the overall sampling rate, we can switch between different algorithms (even choosing to turn off sampling), since as long as each output stratum is a simple random sample of the desired size, the overall output will be a stratified random sample.

This ability to choose to sample the different relations at the stratum level is important in order to minimize the sample size and possible only in stratified random sampling as opposed to simple random sampling. This guides us towards the StratJoin_Overall algorithm provided below.

5. STRATJOIN_OVERALL

StratJoin_Overall provides a stratified random sample of a join while sampling both relations and minimizes the sample size by handling sample inflation. Using the fact that the sampling rate of a stratum a for the relation R_1 can be given by $f \times m_2(a)$, we present a new algorithm that takes the cardinalities of individual strata and the sampling rate into consideration to decide the best sampling and join strategies in order to minimize the sample size. **This simple algorithm provides the best of both worlds – it reduces the sample size when sampling is beneficial and avoids sampling when it is counter-productive.**

Algorithm:
1. For both relations, sample a stratum a with replacement if $f \times m_{other}(a) < 1$. Otherwise, use the entire stratum in the sample.
2. For a stratum, depending on whether none of the relations, one of the relations, or both of the relations are sampled, join the samples as explained below.

(a) If neither relation is sampled, choose random tuples with replacement from both the strata and join them till the desired sampling rate is achieved.

(b) If a single relation is sampled, while tuples from the sampled stratum are streaming by, join them with a randomly chosen tuple from the non-sampled stratum. This builds upon the StratJoin_nn algorithm, which samples a single relation (Section 4.2).

(c) If both strata are sampled, use *Mini-Join* operator between the two samples. This builds upon the StratJoin_Both algorithm, which samples both relations (Section 4.3).

Figure 3 helps explain the algorithm using a walkthrough of a stratum participating in the join.

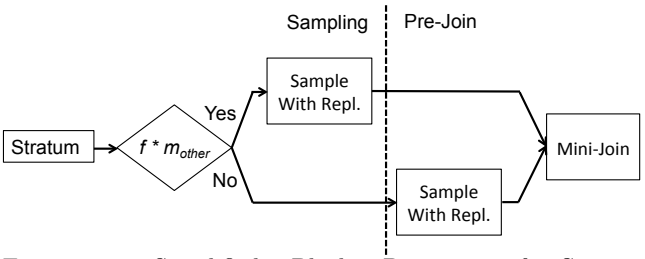

Figure 3: Simplified Block Diagram of StratJoin_Overall

Complexity:
In the worst case, $f \times m_1$ and $f \times m_2$ will be greater than 1 for every stratum in both the relations and as a result, none of the strata will be sampled with the resultant sample size being $|R_1| + |R_2|$. In the best case, where all the strata from both the relations can be sampled, the complexity will be $O(f \times |R_1 \bowtie R_2|)$. The join complexity in each case will be $f \times |R_1 \bowtie R_2|$.

Theorem 7 StratJoin_Overall results in a with-replacement *stratified* random sample of $R_1 \bowtie R_2$, while sampling *both* relations and *minimizes* the sample size.

Proof of Correctness:
Satisfaction of $StRS_1$: The algorithm ensures that in each of the three scenarios, the output stratum is of correct size.
Satisfaction of $StRS_2$ & $StRS_3$: For case (a), since we are repeatedly choosing samples with replacement from both sides to join, it is clear that the output stratum will be a sample with replacement from the join of size $m_1(a) \times m_2(a)$. For case (b), the proof of correctness at the level of stratum for StratJoin_NN will be applicable. Similarly, for case (c), the proof of correctness at the level of stratum for StratJoin_Both will be applicable. It is clear that using smaller size samples will induce correlation between output tuples.

5.1 Handling Sample Inflation in StratJoin_Overall

i	m_1^i	m_2^i	No Sampling	StratJoin_Overall	SRS_Both	Stream-Sample
1	1000	5	1005	500+5 =505	500+**500** =1000	500+5 =505
2	1000	15	1015	1000+15 =1015	**1500**+**1500** =**3000**	**1500**+15 =**1515**
3	5	1000	1005	5+500 =505	**500**+500 =1000	**500**+1000 =**1500**
4	15	1000	1015	15+1000 =1015	**1500**+**1500** =**3000**	**1500**+1000 =**2500**

Table 2: StratJoin_Overall tackling Sample Inflation, $f = 0.1$

Table 2 demonstrates the benefits of StratJoin_Overall and its handling of the sample inflation issue. Column i denotes the stratum index and m_1^i and m_2^i denote the strata counts for stratum i in R_1 and R_2, respectively. Other columns represent the corresponding sampling techniques. In those columns, the values being added are the corresponding sample strata counts as a result of sampling m_1^i and m_2^i respectively. The values in the columns for Srs_Both and *Stream-Sample* denote the average output counts (Section 3.3.4). A sampling rate of 0.1 is used. Instances where the sample stratum size is greater than the relation stratum size have been made bold for emphasis.

In this example, we can see that sampling is beneficial only when the stratum size in the other relation is 5, which is true only in the case of m_1^1 and m_2^3. Since Srs_Both and *Stream-Sample* do not check if the sample stratum size will be bigger than the relation stratum size, we can see that sampling inflation results in the sample stratum size exceeding the relation stratum size for them for numerous strata.

StratJoin_Overall avoids sampling inflation by checking for feasibility while constructing the sample. Therefore, it always decides on the minimum possible sample size for each stratum, thereby, reducing the overall sample size.

5.2 Reduction in Sample Size

We provide below the reduction in the number of tuples that need to participate in the join.

$$\sum_{i=1}^{2} \sum_{j=1}^{Strata\ Count} max(m_i^j - f \times m_1^j \times m_2^j, 0) \qquad (1)$$

Whenever sampling is feasible, it can result in large savings if the size of the stratum being sampled is large. We can also see that when sampling is feasible, the fraction of tuples saved can be given by $1 - f \times m_{other}$.

5.3 Comparison with Stream-Sample

Stream-Sample is more efficient compared with the then previously available technique of *Strategy Olken-Sample* and has fewer restrictions regarding availability of indexes [6]. One of the major goals of sampling is being able to process smaller sized data. However, as we have seen, due to sample inflation the sample sizes using *Stream-Sample* can get exceedingly large, thereby, negating the benefits of sampling. We will see later in Figure 9, that bigger sample size might result in having to read the samples from the disk – resulting in a *substantial* increase in execution time.

Therefore, knowing when to sample and when not to is *critically* important. Our heuristics result in the sample size never exceeding the size of the relations, thereby, making join sampling a feasible proposition. Further, we enable sampling from *both* sides of a join. We reduce constraints on sampling from the sizes of relations to those of strata. As Table 2 shows, sampling can be greatly beneficial when a stratum size is large and the corresponding $f \times m_{other}$ is low (m_1^1 and m_2^3).

5.4 Sample Size Optimality

As we have seen, with a sampling rate of f and sizes of $m_1(a)$ and $m_2(a)$ for the stratum a, the output needs to have $f \times m_1(a) \times m_2(a)$ tuples. However, if $f \times m_1(a) \times m_2(a)$ is greater than $m_1(a)$, it is clearly counter-productive to sample the stratum a of R_1. Therefore, in order to demonstrate minimization of sample size, we need to show that it is not possible to have a sample of size smaller than $f \times m_1(a) \times m_2(a)$ when $f \times m_1(a) \times m_2(a) < m_1(a)$ in order to provide a *random* sample. As we have seen, having a sample of size smaller than $f \times m_1(a) \times m_2(a)$ will result in some tuples having to be reused, introducing correlation between output tuples – voiding randomness.

5.5 Usefulness of Non-sampling

Joining random tuples from relations till the required sampling rate is met can be a better alternative to both *Stream-Sample* and *Naive-Sample* as given by Chaudhari et al. [6]. It is easy to see that in *Naive-Sample*, the entire join is materialized and then sample is generated – allowing *Stream-Sample* to perform join faster. As we have seen in Sections 4.4 and 5.1, *Stream-Sample* can result in the sample size being multiple orders of magnitude larger than the relation size and thus *not sampling* can actually result in the sample size being comparatively smaller.

6. MAXIMIZING JOIN RANDOMNESS

Although obtaining a stratified random sample of join is more feasible than obtaining a simple random sample using our techniques, join sampling is inherently expensive due to sample inflation. In the cases where random samples are preferred but not required, we present sample allocation techniques to increase the randomness of the sample given the constraint of fixed sample size. In order to do so, we introduce a novel concept of *join randomness*. It not only helps quantify the randomness of the join but also provides basis for maximizing it. As mentioned before, non-random sampling in joins has been researched in the context of aggregation queries. This is the first work to look at the degree of randomness in non-random samples.

6.1 Uniformity Confidence

Uniformity Confidence (UC) has been used in order to quantify randomness of a sample [4] and can be given by

$$UC = \frac{\begin{array}{c}The\ number\ of\ different\ samples\ of\ the\\ same\ size\ possible\ with\ the\ algorithm\end{array}}{\begin{array}{c}The\ number\ of\ different\ samples\ of\\ the\ same\ size\ possible\ statistically\end{array}} \times 100$$

We use this metric for two reasons. First, we wish to provide techniques to increase the randomness of the sample. Second, we wish to present the value of the UC metric to the user so that he can increase the sample size if needed. In Sections 6.2 and 6.3, we consider strategies to increase the number of possible samples (numerator of UC), since the number of maximum samples possible statistically (denominator of UC) is independent of any allocation strategy. We first find the optimal allocation in the case of multiple relations joining on a single stratum and then extend it towards the case of multiple relations joining on multiple strata, in the context of sampling without replacement. Section 7.3 shows that our heuristics yield high quality approximations to the best possible allocation found using brute force. We prove the validity of our techniques theoretically in Section 10.2.

6.2 Maximizing Randomness for Single Stratum

Consider having to allocate k^j tuples amongst relations $R_1, R_2 \ldots R_z$, each having a single stratum in order to increase the number of possible samples. Thus, $\sum_{i=1}^{z} mm_i^j = k^j$.

Symbol	Explanation
z	Number of Relations
n	Number of Strata
k	Total Sample Size
k^j	Number of Samples Allocated for Stratum j for all Relations

Table 3: List of new Notations used in Section 6

The number of possible samples can be given by $\prod_{i=1}^{z} C_{mm_i^j}^{m_i^j}$ We use the following allocation strategy to maximize it.

$$mm_i^j = round(\frac{k^j \times m_i^j}{\sum_{i=1}^{z} m_i^j}). \quad (2)$$

Since we are considering a single stratum, j is constrained to be 1. The resultant allocation has been shown by our experiments in Section 7.3.1 to have low error with the maximum difference of 2 from the optimal allocation found by searching through all possible allocations. Section 10.2.1 presents its proof.

6.3 Maximizing Randomness for Multiple Strata

We now provide the technique for allocating a given sample size amongst the different strata. After allocating samples to a given stratum, the samples will be divided amongst the same stratum of different relations using Equation 2.

Let k^1, k^2, $\ldots k^n$ samples be allocated amongst the different strata with $\sum_{j=1}^{n} k^j = k$. Further, an allocation k^j is divided amongst the different relations R_1, R_2 $\ldots R_z$ as mm_1^j, mm_2^j $\ldots mm_z^j$ respectively, so that $k^j = \sum_{i=1}^{z} mm_i^j$. Given these preliminaries, we aim to increase the number of possible samples given by $\prod_{j=1}^{n}\prod_{i=1}^{z} C_{mm_i^j}^{m_i^j}$. We do so using the following allocation strategy.

$$k^j = round(\frac{k \times \sum_{i=1}^{z} m_i^j}{\sum_{j=1}^{n}\sum_{i=1}^{z} m_i^j}) \quad (3)$$

This result has been verified in Section 7.3.2 to provide allocation close to perfect allocation with the maximum difference of 1. Section 10.2.2 presents its proof. It should be noted that only those strata having non-zero tuples in all relations should be considered.

6.4 Discussion

Using Lagrange multipliers is a popular method in the sampling community to find approximate optimal strata allocation in closed-form under space or cost constraints [8,19–21,26]. This includes important allocation topics of Neyman allocation and cost-based allocation [21] for finding strata sizes in stratified sampling. We build upon this framework to derive the above optimal strata allocation techniques (proofs in Section 10.2). Sample sizes suggested by these closed-form allocation strategies are close to the optimal solution in practice – with a maximum difference of 2 in the case of single stratum allocation and 1 for multiple strata allocation.

7. EXPERIMENTAL EVALUATION

The algorithms were implemented using Java 8 (Sections 3, 4, 5 and 6) and Matlab R2014b (Section 6). Experiments were performed on an Ubuntu Linux 14.04.1 LTS system with a 24-core 2.4GHz Intel Xeon CPU, 256GB DDR3 @ 1866 MHz memory, and a 500GB @ 7200 RPM disk.

In Sections 7.1 and 7.2, we study different aspects of StratJoin_Overall since it is our central algorithm for obtaining perfectly random samples. We also compare it with the state of the art, *Stream-Sample*. We use the TPC-DS benchmark and a similar experimental setup to that used by Chaudhari et al. [6]. The results have conformed to our prior expectations. Section 7.3 studies effectiveness of our sample allocation techniques to maximize join randomness.

7.1 TPC-DS Based Evaluation of StratJoin_Overall

StratJoin_Overall is designed for the case where multiple tuples from a relation can join with multiple tuples from the other. **8** out of **99** queries in the TPC-DS benchmark were found to be well-suited for this use case. We ran these queries on the TPC-DS dataset, having sizes ranging from 2 GB to 10 GB, for different sampling rates, using Postgres 9.3. Different metrics were used to study various features of StratJoin_Overall. Unless otherwise stated, each measurement represents the *average* across queries, unless otherwise stated.

7.1.1 *Ratio of* StratJoin_Overall *Sample Size and Relation Size*

As the sampling rate increases, fewer strata can be sampled due to the increasing sample size, resulting in the benefits of StratJoin_Overall decreasing (Figure 4). As expected, as the sampling fraction approaches 1, the ratio of StratJoin_Overall sample size and relation size approaches 1.

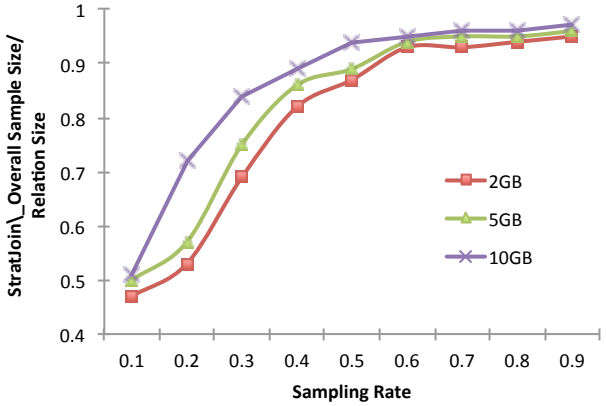

Figure 4: As sampling becomes less feasible with increasing rate, StratJoin_Overall sample size approaches the relation size.

7.1.2 *Ratio of Sample Sizes of* Stream-Sample *and* StratJoin_Overall

StratJoin_Overall explicitly samples both sides of the join while *Stream-Sample* samples a single relation. At low sampling rates, it is more feasible to sample from both the relations. Therefore, the ratio of the sample sizes is more pronounced. The benefit of sampling decreases with increasing rate (Figure 5) as the stratum sample sizes approach the stratum sizes. However, *Stream-Sample* continues sampling even when infeasible (around sampling rate of 0.4) resulting in *Stream-Sample* sample size being greater than the StratJoin_Overall sample size.

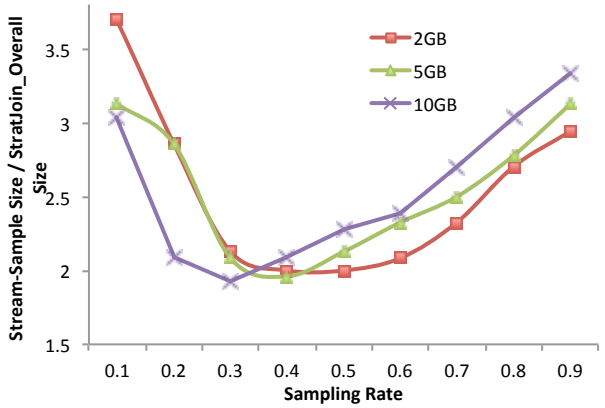

Figure 5: Sample size of *Stream-Sample* is consistently greater than StratJoin_Overall sample size.

7.1.3 *Ratio of Sample Sizes of* Stream-Sample *and* StratJoin_Overall *During Infeasibility*

Figure 5 has demonstrated the increase in sample size of *Stream-Sample* due to sampling during infeasibility caused by sample inflation. Figure 6 shows the benefit of our gating condition which detects this infeasibility in greater detail. (The median sample size ratio is provided since the average ratio is greater by multiple orders of magnitude due to a couple of outlying query results.) As expected, the ratio increases approximately linearly with the sampling rate, which causes the sample size to increase linearly as well.

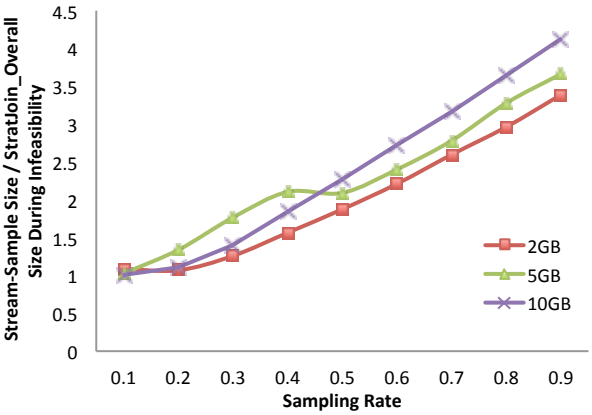

Figure 6: With increasing sampling rate, size of infeasible samples using *Stream-Sample* increases.

7.1.4 *Fraction of Strata Sampled*

We have seen that the gating condition for sampling feasibility employed by StratJoin_Overall, which although can be restrictive, always results in the smallest possible sample size. Here, we look at its applicability by finding out the fraction of strata where sampling is feasible (Figure 7). There exist a couple of queries with large strata sizes where none of the strata can be sampled resulting in the fraction being around 0.6, even for sampling rate of 0.1. As is obvious, we can see that with increasing sampling rate, the number of strata that can be sampled decreases, thereby, reducing the applicability of sampling.

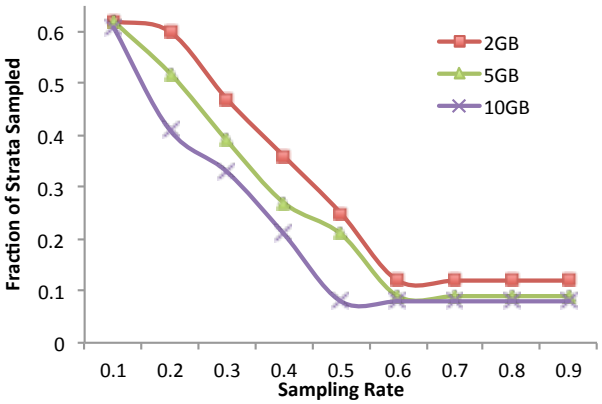

Figure 7: Fraction of strata sampled decreases with increasing sampling rate.

7.1.5 *Fraction of Queries Partially Sampled*

Here we look at the queries having at least a single stratum being sampled. We see that with increasing sampling rate, the fraction of queries having at least 1 stratum that can be sampled starts decreasing (Figure 8). However, even for high sampling rates, most queries contain some strata that can be sampled. The line plot for the *2 GB* case is not visible since it overlaps with that of the *5 GB* one.

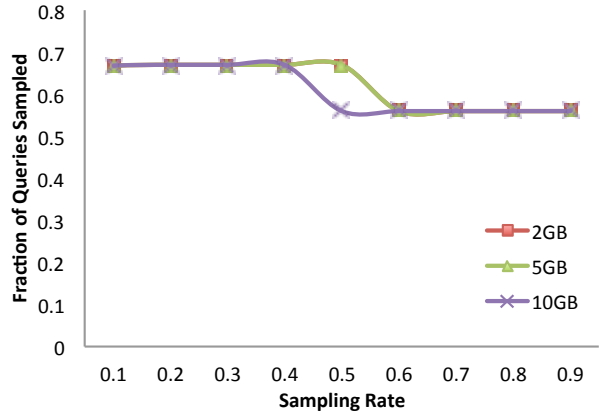

Figure 8: Fraction of queries that are sampled decreases with increasing sampling rate.

7.2 In-Depth Comparison with Stream-Sample

In this section, we compare execution times of StratJoin_Overall and *Stream-Sample*. We use a similar experimental setup to that used by them [6]. Four tables were generated with 100K tuples each. The join column in each table had strata counts derived from a zipfian distribution. The parameter z of the zipfian distributin varied from 0 to 3. Additional four tables with 1M tuples were generated similarly. Each row in the table consists of three columns – RID (integer), JoinKey (integer), and Padding (integer). LHS refers to the relation with 100K tuples while RHS consists of 1M tuples.

7.2.1 *Varying Sampling Rate*

The join complexity of StratJoin_Overall and *Stream-Sample* is $f \times |R_1 \bowtie R_2|$. Therefore, StratJoin_Overall

and *Stream-Sample* require nearly equal time when the samples can be stored in memory, z = 0 for both relations (Figure 9). The execution duration increases nearly proportionally with the sampling rate. However, due to sample inflation, the sample size in *Stream-Sample* can get exceedingly large (f = 0.01 and 0.1, z = 3 for both relations). This results in the sample needing to be written to disk for *Stream-Sample* after sampling and then read from the disk during the join causing slowdown by an order of magnitude, illustrating the need for the gating condition.

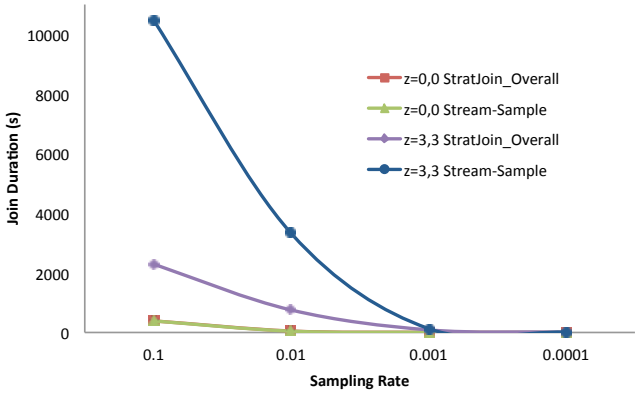

Figure 9: Sample Inflation causes samples to be read from disk increasing the execution time for *Stream-Sample*.

7.2.2 *Varying Skew*

With increased skew, the join size and thereby join duration increases (Figure 10). With the RHS z values of 2 and 3, and LHS z value of 3, samples need to read from disk for *Stream-Sample* resulting again in significant increase in execution duration.

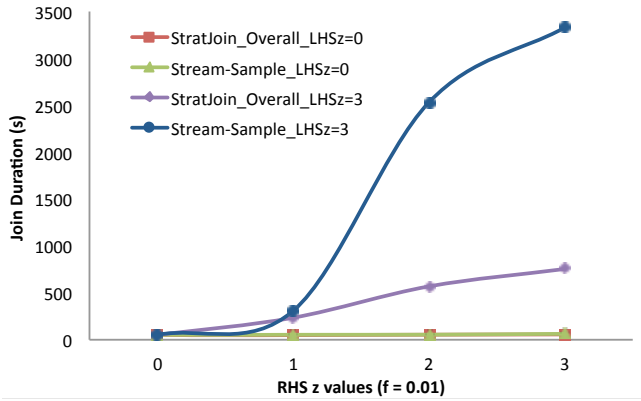

Figure 10: With increased skew, the join size increases causing samples to be read from disk for *Stream-Sample*.

7.3 Maximizing Join Randomness

In this section, we study the effectiveness of the sample allocation techniques. We first validate our allocation strategy for a single stratum and then for multiple strata. Since finding the optimal solution by considering all possible combinations is computationally expensive, we have restricted the size of our datasets.

The effectiveness of our sample allocation techniques was studied by comparing their results with those of the best possible allocation found by searching all possible allocations and providing *Mean-Squared Error*, *Mean-Squared Relative Error*[1], and maximum difference.

7.3.1 *Single Stratum Partitioning*

mm_1^1	$k = 6$	$k = 12$	$k = 18$
1	155040	1679600	11400
2	**218025**	8314020	218025
3	136800	20155200	1860480
4	39900	**26453700**	8139600
5	5040	19535040	19535040
6	210	8139600	**26453700**
7	-	1860480	20155200
8	-	218025	8314020
9	-	11400	1679600
10	-	190	125970

Table 4: Number of Possible Samples

We first provide the allocation for an illustrative use case. It enables an intuitive understanding of the behavior of the numerator of UC. Table 4 shows the number of possible samples when k^1 tuples need to be assigned between two relations having a single stratum. Here, $m_1^1 = 10$ and $m_2^1 = 20$. It can be easily verified that Equation 2 results in the maximum number of possible samples.

Approximation Error:

Population	MSE	MSRE	Maximum Difference
150	0.1109	0.0069	1
200	0.1197	0.0094	2
300	0.1377	0.0135	2
400	0.1506	0.0162	2
500	0.1598	0.0180	2
600	0.1666	0.0196	2
700	0.1720	0.0208	2
800	0.1761	0.0218	2
900	0.1794	0.0226	2
1000	0.1821	0.0232	2

Table 5: Error in Single Stratum Partitioning

However, our technique might not result in ideal allocation for all possible sample size configurations. Hence, here we study their accuracy. We used 3 relations. The population, varying from 150 to 1000, was divided in all combinations between the 3 relations such that each relation at least had 5 tuples. The sample size was varied from 15 to 100. Table 5 shows that all error metrics are low – validating the effectiveness of our formula in practice for the single stratum use case.

7.3.2 *Multiple Strata Partitioning*

We proceed with an illustrative use case as before. Figure 11 shows the number of possible samples for two relations having two strata each, with the configuration of $m_1^1 = 10$, $m_2^1 = 20$, $m_1^2 = 5$, $m_2^2 = 15$, and $k = 20$. We can see that our heuristic resulted in the maximum number of samples as $k^1 = round(\frac{20\times(10+20)}{10+20+5+15}) = 12$.

[1]defined as $\frac{\sum_{i=1}^{n}(\frac{actual(i)-predicted(i)}{actual(i)})^2}{n}$

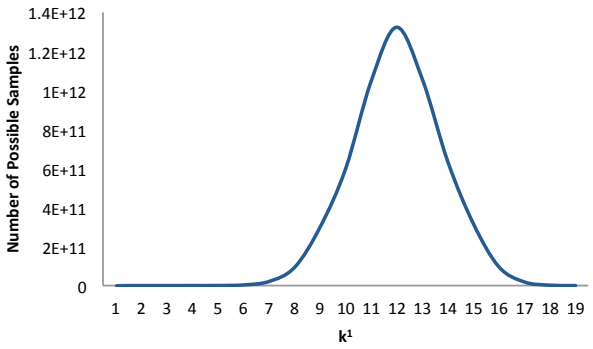

Figure 11: Number of Possible Samples for Different k^1

Approximation Error:

Population	MSE	MSRE	Maximum Difference
40	0.1439	0.0183	1
45	0.1714	0.0221	1
50	0.1097	0.0081	1
55	0.1087	0.0128	1

Table 6: Error in Multiple Strata Partitioning

We again studied the effectiveness of our strategy using varying population and sample sizes. We used 3 relations having 3 strata each with the minimum size of a stratum being 3. The total sample size was varied from 27 to 30 and the population from 40 to 55. The strata sizes were varied to cover every possible combination. Table 6 shows that all error metrics are again low. As mentioned earlier, the population size was restricted due to the computational complexity of finding the optimal allocation using brute force.

7.3.3 Heuristic Discussion

Our heuristics maximize the number of samples in the continuous domain whereas the allocation exists in the discrete domain, in order to derive the allocation strategies. Thus, the heuristics provide optimal solution when the rounding function does not have any effect. The approximation errors being low further affirms their usefulness in practice.

8. CONCLUSION & FUTURE WORK

Sampling using both sides of a join is an important area of work. In this paper, we have presented techniques to sample joins, in the context of perfect as well as maximum randomness – illustrating the benefits and drawbacks in doing so. In the context of perfect randomness, we have provided algorithms for obtaining stratified random samples of a join by sampling a single relation or both relations of the join. We show how the sample inflation issue can be handled using stratified random sampling resulting in minimization of the sample size. We provided novel techniques for increasing randomness of joins when randomness is preferable but not mandatory. Our experiments confirmed that the benefits of sampling reduce with increasing sampling rate due to increasing sample size. Further, we saw that not handling sample inflation results in the sample size far exceeding the relation sizes. In some cases, this had a detrimental effect of sampled tuples being needed to be read from the disk – further compounding the ill-effects of sample inflation. We also saw that our techniques to maximize the join randomness were effective over varying population and sample sizes.

Going forward, several interesting avenues of research are enabled by our findings. First, there are opportunities into improving the efficiency of our algorithms. A second appealing area of investigation is to compare and contrast our work in the context of online aggregation research. Specifically, it would be interesting to confirm if our techniques on maximizing randomness of the sample can benefit aggregate estimation using correlated samples. We would also like to investigate the applicability of a search-based solution for maximizing randomness [27]. Third, we plan to look at the performance aspects of our approach by considering hashing, or auxiliary data structures such as indexes (inspired by outlier indexing [5]) to store exact or approximate distribution information. Fourth, we would also like to develop algorithms that consider different levels of availabilities of indexes and statistics since high degree of presence of indexes in numerous use cases such as OLAP is uncommon. Finally, we hope that our work enables applications of join sampling that were so far infeasible, especially in the areas of approximate data mining, exploratory visualization, and distributed query processing, and that the necessity or lack thereof of randomness is taken into consideration while choosing the algorithms.

9. REFERENCES

[1] S. Acharya et al. Join Synopses for Approximate Query Answering. *SIGMOD*, 1999.

[2] S. Acharya et al. The Aqua Approximate Query Answering System. *SIGMOD*, 1999.

[3] S. Agarwal, B. Mozafari, A. Panda, et al. BlinkDB: Queries With Bounded Errors and Bounded Response Times on Very Large Data. *EuroSys*, 2013.

[4] M. Al-Kateb, B. S. Lee, et al. Adaptive-size Reservoir Sampling Over Data Streams. *SSDBM*, 2007.

[5] S. Chaudhuri, G. Das, et al. Overcoming limitations of sampling for aggregation queries. *ICDE*, 2001.

[6] S. Chaudhuri, R. Motwani, and V. Narasayya. On random sampling over joins. *SIGMOD*, 1999.

[7] W. G. Cochran. *Sampling Techniques*. John Wiley & Sons, 2007.

[8] J. A. Dıaz-Garcıa et al. Optimum allocation in stratified surveys. *Comunicaciones del Cimat*, 2005.

[9] A. Dobra, C. Jermaine, et al. Turbo-charging Estimate Convergence in DBO. *VLDB*, 2009.

[10] P. B. Gibbons, V. Poosala, S. Acharya, Y. Bartal, Y. Matias, et al. AQUA: System and Techniques for Approximate Query Answering. *Bell Labs TR*, 1998.

[11] P. J. Haas. Large-sample and Deterministic Confidence Intervals for Online Aggregation. *SSDBM*, 1997.

[12] P. J. Haas et al. Selectivity and Cost Estimation for Joins Based on Random Sampling. *JCSS*, 1996.

[13] P. J. Haas and P. J. Haas. Hoeffding Inequalities for Join-selectivity Estimation and Online Aggregation. *Research Report*, 1996.

[14] P. J. Haas and J. M. Hellerstein. Ripple Joins for Online Aggregation. *SIGMOD*, 1999.

[15] J. Hellerstein et al. Online Aggregation. *SIGMOD*, 1997.

[16] C. Jermaine, A. Dobra, S. Arumugam, S. Joshi, and A. Pol. The Sort-Merge-Shrink Join. *TODS*, 2006.
[17] C. Jermaine, A. Dobra, et al. A Disk-based Join with Probabilistic Guarantees. *SIGMOD*, 2005.
[18] C. Jermaine et al. Scalable Approximate Query Processing with the DBO Engine. *TODS*, 2008.
[19] J. B. Kadane. Optimal dynamic sample allocation among strata. *Department of Statistics*, 2005.
[20] K. Kitikidou. Optimizing forest sampling by using lagrange multipliers. *AJOR*, 2012.
[21] S. L. Lohr. *Sampling: Design and Analysis*. Cengage Learning, 2010.
[22] G. Luo, C. J. Ellmann, P. J. Haas, et al. A Scalable Hash Ripple Join Algorithm. *SIGMOD*, 2002.
[23] S. Nirkhiwale, A. Dobra, and C. Jermaine. A Sampling Algebra for Aggregate Estimation. *VLDB*, 2013.
[24] F. Olken. *Random Sampling from Databases*. PhD thesis, University of California, 1993.
[25] L. Sidirourgos, M. L. Kersten, and P. A. Boncz. Sciborq: Scientific data management with bounds on runtime and quality. *CIDR*, 2011.
[26] P. V. Sukhatme. *Sampling theory of surveys with applications*. ISAS; New Delhi, 1957.
[27] B. W. Wah and Z. Wu. The Theory of Discrete Lagrange Multipliers for Nonlinear Discrete Optimization. *CP*, 1999.

10. APPENDIX

10.1 Proof of Correctness for *Stream-Sample*

While Chaudhari et. al. have provided the *Stream-Sample* algorithm in [6], they do not show its correctness, which we do. This is crucial towards understanding and designing join sampling algorithms. We prove that this algorithm gives a with replacement simple random sample of the join by showing that it satisfies SRS_1 and SRS_2.

Satisfaction of SRS_1: Consider a tuple t present in the output at the index i, which is the result of the join of t_1 at index i in S_1 and t_2 chosen randomly from R_2.
$P(t_1 \in S_1 \text{ at index } i) = \frac{m_2(t1.A)}{|R_1 \bowtie R_2|}$
$P(t_2 \text{ being chosen}) = \frac{1}{m_2(t1.A)}$
$P(t \in S_{12} \text{ at index } i) = \frac{m_2(t1.A)}{|R_1 \bowtie R_2|} \times \frac{1}{m_2(t1.A)} = \frac{1}{|R_1 \bowtie R_2|}$
We can thus see that the probability of any tuple being present in the output at any index is the same and as expected equal to $\frac{1}{|R_1 \bowtie R_2|}$.

Satisfaction of SRS_2: The tuples in S_1 are selected independently of each other. Tuples are selected from R_2 independently as well. Therefore, there is no correlation between two tuples in the output.

10.2 Proofs for Maximizing Randomness

10.2.1 Maximizing Randomness for Single Stratum

The number of possible samples under sampling without replacement for an equi-join can be given by

$$f(mm_1, mm_2, \dots mm_z) = \prod_{i=1}^{z} C_{mm_i}^{m_i} = \prod_{i=1}^{z} \frac{m_i!}{mm_i!(m_i - mm_i)!} \tag{4}$$

The notation used in Tables 1 and 3 is used in this subsection and the next. Note that in that notation the superscript denotes the stratum index and the subscript denotes the relation index. The superscript has been omitted from the usual notation of m_i^1 and mm_i^1 since a single stratum is under consideration. The constraint of fixed sample size can be given by

$$g(mm_1, mm_2, \dots mm_z) = \sum_{i=1}^{z} mm_i = k^1 \tag{5}$$

k^1 denotes the number of samples being allocated to the first stratum and g denotes the constraint function. Lagrange multipliers gives us a system of z equations. The i^{th} ($i \in [1, z]$) equation is obtained by taking partial derivative of equation 4 with respect to mm_i

$$f_{mm_i} = \Big(\prod_{j=1 \& j \neq i}^{z} \frac{m_j!}{mm_j! \times (m_j - mm_j)!}\Big) \times m_i! \times \frac{\partial}{\partial mm_i}\Big(\frac{1}{mm_i! \times (m_i - mm_i)!}\Big) \tag{6}$$

Using the basic Lagrange multipliers equation of $f_{mm_i} = \lambda \times g_{mm_i}$ and equation 5, we get

$$f_{mm_i} = \lambda \times g_{mm_i} = \lambda \times \frac{\partial(\sum_{i=1}^{z} mm_i)}{\partial mm_i} = \lambda \tag{7}$$

Using the Gamma function to represent factorial, its derivative can be given by

$$(n!)' = \Gamma'(n+1) = n! \times (-\gamma + \sum_{a=1}^{n} \frac{1}{a}) \tag{8}$$

where γ is Euler-Mascheroni constant. Using equation 8, we can simplify $\frac{\partial}{\partial mm_i}(\frac{1}{mm_i! \times (m_i - mm_i)!})$,

$$- \frac{(mm_i! \times \frac{\partial((m_i - mm_i)!)}{\partial mm_i}) + ((m_i - mm_i)! \times \frac{\partial(mm_i!)}{\partial mm_i})}{(mm_i! \times (m_i - mm_i)!)^2}$$

$$= -\frac{(mm_i! \times (m_i - mm_i)! \times (-\gamma + \sum_{a=1}^{m_i - mm_i} \frac{1}{a}) \times (-1))}{(mm_i! \times (m_i - mm_i)!)^2}$$

$$- \frac{((m_i - mm_i)! \times mm_i! \times (-\gamma + \sum_{a=1}^{mm_i} \frac{1}{a}))}{(mm_i! \times (m_i - mm_i)!)^2}$$

$$= \frac{-(mm_i! \times (m_i - mm_i)!) \times (\gamma - \sum_{a=1}^{m_i - mm_i} \frac{1}{a} - \gamma + \sum_{a=1}^{mm_i} \frac{1}{a})}{(mm_i! \times (m_i - mm_i)!)^2}$$

$$= \frac{-(\sum_{a=1}^{mm_i} \frac{1}{a} - \sum_{a=1}^{m_i - mm_i} \frac{1}{a})}{mm_i! \times (m_i - mm_i)!} \tag{9}$$

Plugging above into equation 6 and using equation 7,

$$\lambda = \Big(\prod_{j=1 \& j \neq i}^{z} \frac{m_j!}{mm_j! \times (m_j - mm_j)!}\Big) \times m_i! \times \frac{(-(\sum_{a=1}^{mm_i} \frac{1}{a} - \sum_{a=1}^{m_i - mm_i} \frac{1}{a}))}{mm_i! \times (m_i - mm_i)!}$$

$$= -\Big(\prod_{j=1}^{z} \frac{m_j!}{mm_j! \times (m_j - mm_j)!}\Big) \times \Big(\sum_{a=1}^{mm_i} \frac{1}{a} - \sum_{a=1}^{m_i - mm_i} \frac{1}{a}\Big) \tag{10}$$

The sum of harmonic series is approximated by

$$\sum_{a=1}^{mm_i} \frac{1}{a} = ln(mm_i) + \gamma + \frac{1}{2 \times mm_i} - \frac{1}{12 \times mm_i^2} + \frac{1}{120 \times mm_i^4}$$
$$\approx ln(mm_i) + \gamma \quad (11)$$

for large mm_i where γ is the Euler-Mascheroni constant. Using equations 10 and 11,

$$-\lambda = (\prod_{j=1}^{z} \frac{m_j!}{mm_j! \times (m_j - mm_j)!}) \times$$
$$(ln(mm_i) + \gamma - ln(m_i - mm_i) - \gamma)$$
$$= (\prod_{j=1}^{z} \frac{m_j!}{mm_j! \times (m_j - mm_j)!}) \times ln(\frac{mm_i}{m_i - mm_i}) \quad (12)$$

Note that the term $(\prod_{j=1}^{z} \frac{m_j!}{mm_j! \times (m_j - mm_j)!})$ is common for all $i \in [1, z]$ equations and we denote it by C. The above system of equations can thus be simplified as

$$\frac{-\lambda}{C} = ln(\frac{mm_i}{m_i - mm_i}) \quad (13)$$

Consider the allocation mm_p for another relation p and its corresponding equation given by

$$\frac{-\lambda}{C} = ln(\frac{mm_p}{m_p - mm_p}) \quad (14)$$

From the above two equations,

$$\frac{mm_i}{m_i - mm_i} = \frac{mm_p}{m_p - mm_p} = A \quad (15)$$

for a constant A. Therefore,

$$mm_i = \frac{A \times m_i}{1 + A} \quad (16)$$

Using equations 5 and 16

$$k^1 = \sum_{i=1}^{z} mm_i = \sum_{i=1}^{z} \frac{A \times m_i}{1 + A} = \frac{A}{1 + A} \sum_{i=1}^{z} m_i$$

$$\frac{A}{1 + A} = \frac{k^1}{\sum_{i=1}^{z} m_i} \quad (17)$$

Using equations 16 and 17

$$mm_i = (\frac{A}{1 + A}) \times m_i = \frac{k^1 \times m_i}{\sum_{a=1}^{z} m_a} \quad (18)$$

As we have seen from the experiments in Section 7.3.1, this is close to the optimal allocation found using brute force.

10.2.2 Maximizing Randomness for Multiple Strata

This proof is similar to the one in the previous section. In this case, our goal is to maximize the number of possible samples under sampling without replacement. The constraint of fixed sample size is given by

$$k = \sum_{j=1}^{n} k^j = g(k^1, k^2 \dots k^n) \quad (19)$$

where g represents the constraint function. Also,

$$k^j = \sum_{i=1}^{z} mm_i^j \quad (20)$$

The number of possible samples can be given by

$$f(k^1, k^2 \dots k^n) = \prod_{j=1}^{n} \prod_{i=1}^{z} C_{mm_i^j}^{m_i^j}$$
$$= \prod_{j=1}^{n} \prod_{i=1}^{z} \frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!}$$
$$= \prod_{j=1}^{n} ((\prod_{i=1}^{z-1} \frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!}) \times \frac{m_z^j!}{mm_z^j! \times (m_z^j - mm_z^j)!})$$
$$= \prod_{j=1}^{n} ((\prod_{i=1}^{z-1} \frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!}) \times$$
$$\frac{m_z^j!}{(k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))!}) \quad (21)$$

Equation 18 in the previous section provides the optimal allocation amongst the relations given allocation to a stratum. It is easy to see that once we obtain an optimal allocation, by redistributing the total allocation of a stratum to the different relations using equation 18, the number of possible samples will be greater than or equal to the current number of samples. Therefore, the variables in the above function that needs to be maximized are k^j ($j \in [1, n]$). Using Lagrange multipliers, we get the j^{th} ($j \in [1, n]$) equation by taking partial derivative with respect to k^j as

$$f_{k_j} = \prod_{j=1}^{n} ((\prod_{i=1}^{z-1} \frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!}) \times m_z^j! \times$$
$$\frac{\partial}{\partial k^j} (\frac{1}{(k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))!})) \quad (22)$$

We can simplify $\frac{\partial}{\partial k^j} (\frac{1}{(k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))!})$,

$$\frac{-1}{((k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))!)^2} \times$$
$$(k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (-\gamma + \sum_{l=1}^{k^j - \sum_{i=1}^{z-1} mm_i^j} \frac{1}{a}) \times$$
$$(m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))! - (k^j - \sum_{i=1}^{z-1} mm_i^j)! \times$$
$$(m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))! \times (-\gamma + \sum_{l=1}^{m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j)} \frac{1}{a})$$

$$= \frac{-1 \times (k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))! \times}{((k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))!)^2}$$
$$(-\gamma + \sum_{l=1}^{k^j - \sum_{i=1}^{z-1} mm_i^j} \frac{1}{a} + \gamma - \sum_{l=1}^{m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j)} \frac{1}{a})$$

$$= \frac{-(ln(k^j - \sum_{i=1}^{z-1} mm_i^j) + \gamma - ln(m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j)) - \gamma)}{(k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))!}$$

$$= \frac{ln(\frac{m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j)}{k^j - \sum_{i=1}^{z-1} mm_i^j})}{mm_z^j! \times (m_z^j - mm_z^j)!} \quad (23)$$

Substituting above equation in equation 22,

$$f_{k_j} = (\prod_{j=1}^{n}\prod_{i=1}^{z}\frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!})ln(\frac{m_z^j - k^j + \sum_{i=1}^{z-1} mm_i^j}{k^j - \sum_{i=1}^{z-1} mm_i^j})$$

From the standard Lagrange multipliers equation of $f_{k_j} = \lambda \times g_{k_j}$ and equation 19,

$$f_{k_j} = \lambda \times g_{k_j} = \lambda \times \frac{\partial(\sum_{j=1}^{n} k^j)}{\partial k_j} = \lambda \tag{24}$$

From the above 2 equations,

$$\lambda = (\prod_{j=1}^{n}\prod_{i=1}^{z}\frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!})ln(\frac{m_z^j}{k^j - \sum_{i=1}^{z-1} mm_i^j} - 1)$$

Consider another stratum p and its corresponding equation

$$\lambda = (\prod_{j=1}^{n}\prod_{i=1}^{z}\frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!})ln(\frac{m_z^p}{k^p - \sum_{i=1}^{z-1} mm_i^p} - 1)$$

From the above 2 equations, given a constant A, we get

$$\frac{k^j - \sum_{i=1}^{z-1} mm_i^j}{m_z^j} = \frac{k^p - \sum_{i=1}^{z-1} mm_i^p}{m_z^p} = A \tag{25}$$

Using equation 18 and the above equation,

$$A = \frac{k^j - \sum_{i=1}^{z-1} \frac{k^j \times m_i^j}{\sum_{a=1}^{z} m_a^j}}{m_z^j} = \frac{k^j \times (1 - \sum_{i=1}^{z-1} \frac{m_i^j}{\sum_{a=1}^{z} m_a^j})}{m_z^j}$$
$$= \frac{k^j \times (\sum_{a=1}^{z} m_a^j - \sum_{i=1}^{z-1} m_i^j)}{m_z^j \times \sum_{a=1}^{z} m_a^j} = \frac{k^j}{\sum_{a=1}^{z} m_a^j} \tag{26}$$

$$k^j = A \times \sum_{a=1}^{z} m_a^j \tag{27}$$

Using above equation and equation 19

$$k = \sum_{j=1}^{n}(A \times \sum_{a=1}^{z} m_a^j) \tag{28}$$

$$A = \frac{k}{\sum_{j=1}^{n}\sum_{a=1}^{z} m_a^j} \tag{29}$$

Using equations 27 and 29

$$k^j = \frac{k \times \sum_{i=1}^{z} m_i^j}{\sum_{j=1}^{n}\sum_{i=1}^{z} m_i^j} \tag{30}$$

Section 7.3.2 shows that this approximation gives allocation close to the optimal allocation found using brute force.